

RESEARCH

Parallel Knowledge Embedding with MapReduce on a Multi-core Processor

Miao Fan^{1,2,3*}, Qiang Zhou^{1,2,3}, Thomas Fang Zheng^{1,2,3} and Ralph Grishman^{4,5,6}

*Correspondence:

fanmiao.cslt.thu@gmail.com

¹Center for Speech and Language Technology, Research Institute of Information Technology, Tsinghua University, ROOM 1-303, BLDG FIT, 100084 Beijing, China
Full list of author information is available at the end of the article

Abstract

This article firstly attempts to explore parallel algorithms of learning distributed representations for both entities and relations in large-scale knowledge repositories with *MapReduce* programming model on a multi-core processor.

We accelerate the training progress of a canonical knowledge embedding method, i.e. *translating embedding* (**TransE**) model, by dividing a whole knowledge repository into several balanced subsets, and feeding each subset into an individual core where local embeddings can concurrently run updating during the *Map* phase.

However, it usually suffers from inconsistent low-dimensional vector representations of the same key, which are collected from different *Map* workers, and further leads to conflicts when conducting *Reduce* to merge the various vectors associated with the same key.

Therefore, we try several strategies to acquire the merged embeddings which may not only retain the performance of *entity inference*, *relation prediction*, and even *triplet classification* evaluated by the single-thread **TransE** on several well-known knowledge bases such as Freebase and NELL, but also scale up the learning speed along with the number of cores within a processor.

So far, the empirical studies show that we could achieve comparable results as the single-thread **TransE** performs by the *stochastic gradient descend* (SGD) algorithm, as well as increase the training speed multiple times via adapting the *batch gradient descend* (BGD) algorithm for *MapReduce* paradigm.

Keywords: parallel knowledge embedding; MapReduce; stochastic gradient descend (SGD); batch gradient descend (BGD)

1 Introduction

The emerging trend of learning distributed representations for both entities and relationships using knowledge repositories [1] without extra text has drawn much attention from academia as well as industry. The learnt embeddings can not only perform significant improvements on several subtasks for knowledge population such as *entity inference*, *relation prediction* and *triplet identification*, but also further facilitate other AI applications, such as *Question Answering Systems* and *Automatic Summarization Systems*. Almost all the studies of knowledge embedding leverage existing repositories such as Wordnet^[1] [2], Freebase^[2] [3, 4] and NELL^[3] [5], which contain billions of triplets, and each of them is represented by $\langle \text{head_entity}, \text{relation}, \text{tail_entity} \rangle$ abbreviated as $\langle h, r, t \rangle$. Furthermore, each

^[1]<https://wordnet.princeton.edu/>

^[2]<https://www.freebase.com/>

^[3]<http://rtw.ml.cmu.edu/rtw/>

knowledge base is considered as a huge directed graph where nodes are entities, edges are relations and a triplet $\langle h, r, t \rangle$ implies a structure that there is a relationship r between the entities h and t , pointing from the head entity h to the tail entity t . This hierarchical structure inspired Bordes et al. [6] to interpret relationships as translation vectors operating between the embeddings of two entities. Hence, they have proposed a canonical model named **TransE** (translating embedding) and led the prosperous research of knowledge embedding [7, 8, 9, 10, 11, 12, 13] that focuses on devising models easy to train, containing reduced number of parameters and expected to scale up to very large knowledge bases.

However, to the best of our knowledge, the classical approach **TransE** [6] and its successors [7, 8, 9, 10, 11, 12, 13] conventionally reported the outstanding results of experiments conducted on relatively small datasets such as **WN100K** and **FB150K**^[4], but few of them mentioned the efficiency of training with really large scale datasets. The reason we believe is that all the models including the state-of-the-art, only demonstrated the training algorithms within the programming paradigm of single thread, which extremely seals the potential power of multi-core computing.

Multi-core processors enable the computation conducted simultaneously, but reform the paradigm of programming. An intuitive way of parallel computing is to divide the dataset and to feed each core with each subset to perform individual calculations. Therefore, the key challenges of parallel computing with a multi-core processor are: how to schedule the computation tasks, maintain the algorithm of each core and synchronize the shared memory if necessary. To tackle these issues, Google researchers designed a well-known programming model called *MapReduce* [14] for distributed computing, and it has been widely adopted by both commercial and open-source clusters. *MapReduce* generally provides two interfaces to programmers for adapting their single-thread algorithms:

- *Map* is responsible for processing the input data, and emits a set of *intermediate* key/value pairs based on the algorithm devised by programmers.
- *Reduce* accepts an intermediate key and a set of values for the key in the meanwhile. It uses customized function to merge the values together, and finally produces zero or one output for that key.

During the period of operations, an independent master schedules multiple *Map* and *Reduce* workers to run concurrently and makes parallel computing possible.

So far, several quintessential machine learning models have already been adapted to the framework of *MapReduce* by Chu et al. [15]. However, embedding-based models such as *word embedding* [16] and *knowledge embedding* [1] are different from those traditional learning approaches like *logistic regression*, as the data itself constructs the parameter space of embedding-based models. For example, *knowledge embedding* models acquire the low-dimensional vector representation of each entity and relation within a repository, and if we split that knowledge base and distribute subsets into multiple *Map* workers, the parameter space is divided at the same time as well. It usually leads to inconsistent distributed representations of the same key when conducting *Reduce* with those intermediate key/value pairs emitted by many *Map* workers. The problem is still challenging and has not been properly solved yet.

^[4]We change the original names, **WN** and **FB15K**, of the two datasets to emphasize the sizes of them.

Therefore, this paper firstly attempts to explore parallel algorithms of learning distributed representations for both entities and relations in large-scale knowledge repositories with *MapReduce* programming model on a multi-core processor. We accelerate the training progress of the canonical knowledge embedding method, i.e. *translating embedding* (**TransE**) model [6], by dividing a whole knowledge repository into several balanced subsets, and feeding each subset into an individual core where local embeddings can concurrently run updating during the *Map* phase.

To deal with the merging problem of *Reduce* phase caused by inconsistent low-dimensional vector representations of the same key collected from different *Map* cores, we contribute several strategies to acquire the embeddings which may not only retain the performance of *entity inference*, *relation prediction*, and even *triplet classification* evaluated by the single-thread **TransE** on several well-known knowledge bases such as Freebase and NELL, but also scale up the learning speed along with the number of cores within a processor. So far, the empirical studies show that we could achieve comparable results as the single-thread **TransE** performs by the *stochastic gradient descend* (SGD) algorithm as well as increase the training speed multiple times, via adapting the *batch gradient descend* (BGD) algorithm for *MapReduce* paradigm.

Here we arrange the subsequent content of this article as follows: Section 2 helps us to recap the **TransE** model and its single-thread training algorithm based on *stochastic gradient descend* (SGD). Then, we adapt the single-thread algorithm into the framework of *MapReduce* in Section 3. Specifically, Section 3.1 talks about directly parallelizing the SGD-based algorithm within the *MapReduce* paradigm, and proposes three strategies of reducing the conflicted embeddings. Section 3.2, from another perspective, focuses on concurrently updating the gradients rather than the inconsistent embeddings to avoid conflicts, and suggests adapting the *batch gradient descend* (BGD) algorithm instead into the *MapReduce* paradigm. (TO BE CONTINUED...)

2 Singlethread TransE

Starting from 2011, when Bordes et al. [1] engaged in the representation learning of entities and relationships within knowledge bases, the literatures on knowledge embedding had been increasing, but few of them could support large-scale training within acceptable time due to the high complexity of their models. Until the year of 2013, Bordes et al. [6] devised **TransE**, a canonical model which pursued only one low-dimensional vector representation for each entity and each relationship. The reduced parameters required made it possible to scale up to very large datasets. Moreover, **TransE** learnt those distributed representations with a single-layer energy-based model which was rather easier to train.

2.1 Translating Embedding Model

Inspired by the hierarchical structures of the relationship between two entities shown by Figure 1(a), Bordes et al. interpreted relationships as translations operating on entities and assumed that if a triplet $\langle h, r, t \rangle$ holds, then $\mathbf{h} + \mathbf{r} \approx \mathbf{t}$. The boldface characters indicate k dimensional vector representations (\mathbf{h} , \mathbf{r} and \mathbf{t}) taking values

in \mathbb{R}^k . Equation (1) was derived from that assumption to measure the energy of a triplet $\langle h, r, t \rangle$ as follows,

$$d(\mathbf{h}, \mathbf{r}, \mathbf{t}) = \|\mathbf{h} + \mathbf{r} - \mathbf{t}\|, \quad (1)$$

in which we could take either L_1 or L_2 -norm to measure the dissimilarity. Given a training set Δ of triplets $\langle h, r, t \rangle$ composed of the set of entities E ($h, t \in E$), and the set of relations R ($r \in R$), we expected to discriminate the training triplets from corrupted triplets. One corrupted triplet $\langle h', r, t' \rangle$ from its set Δ' was constructed according to Equation (2) by replacing either the head or tail entity randomly of a ground-truth triplet from the training set:

$$\Delta'_{(h,r,t)} = \{(h', r, t) | h' \in E \ \& \ h' \neq h\} \cup \{(h, r, t') | t' \in E \ \& \ t' \neq t\}. \quad (2)$$

As we favored lower energy of training triplets than the corrupted ones, the objective was naturally to *minimize* the loss of a margin-based ranking function \mathcal{L} over all the training triplets:

$$\mathcal{L} = \sum_{(h,r,t) \in \Delta} \sum_{(h',r,t') \in \Delta'_{(h,r,t)}} [\gamma + d(\mathbf{h}, \mathbf{r}, \mathbf{t}) - d(\mathbf{h}', \mathbf{r}, \mathbf{t}')]_+, \quad (3)$$

where $[x]_+$ was the hinge loss function that only considered the positive part of x ($[x]_+ = \max\{0, x\}$), and γ was that positive margin ($\gamma > 0$).

2.2 SGD-based Singlethread Algorithm

Algorithm 1 The SGD-based Singlethread Learning Algorithm of TransE

Require:

Training set $\Delta = \{(h, r, t)\}$, entity set E , relation set R ; dimension of embeddings d , margin γ , learning rate α for \mathcal{L} , convergence threshold ϵ , maximum epoches n .

```

1: for  $\mathbf{r} \in R$  do
2:    $\mathbf{r} := \text{Uniform}(\frac{-6}{\sqrt{d}}, \frac{6}{\sqrt{d}})$ 
3:    $\mathbf{r} := \frac{\mathbf{r}}{\|\mathbf{r}\|}$ 
4: end for
5:  $i := 0$ 
6: while  $\text{Rel.loss} > \epsilon$  and  $i < n$  do
7:   for  $\mathbf{e} \in E$  do
8:      $\mathbf{e} := \text{Uniform}(\frac{-6}{\sqrt{d}}, \frac{6}{\sqrt{d}})$ 
9:      $\mathbf{e} := \frac{\mathbf{e}}{\|\mathbf{e}\|}$ 
10:  end for
11:  for  $(h, r, t) \in \Delta$  do
12:
13:  end for
14: end while
```

Ensure:

All the embeddings of e and r , where $e \in E$ and $r \in R$.

3 MapReduce TransE

3.1 SGD-based MapReduce Paradigm

3.1.1 SGD-based Map phase

3.1.2 SGD-based Reduce phase

- *Random knowledge embedding*

- Average knowledge embedding
- Mini-loss knowledge embedding

3.2 BGD-based MapReduce Paradigm

3.2.1 BGD-based Map phase

3.2.2 BGD-based Reduce phase

4 Experiment

5 Discussion

6 Conclusion

This article firstly attempts to explore parallel algorithms of learning distributed representations for both entities and relations in large-scale knowledge repositories with *MapReduce* programming model on a multi-core processor.

We accelerate the training progress of a canonical knowledge embedding method, i.e. *translating embedding* (**TransE**) model, by dividing a whole knowledge repository into several balanced subsets, and feeding each subset into an individual core where local embeddings can concurrently run updating during the *Map* phase.

However, it usually suffers from inconsistent low-dimensional vector representations of the same key, which are collected from different *Map* workers, and further leads to conflicts when conducting *Reduce* to merge the various vectors associated with the same key.

Therefore, we try several strategies to acquire the merged embeddings which may not only retain the performance of *entity inference*, *relation prediction*, and even *triplet classification* evaluated by the single-thread **TransE** on several well-known knowledge bases such as Freebase and NELL, but also scale up the learning speed along with the number of cores within a processor.

So far, the empirical studies show that we could achieve comparable results as the single-thread **TransE** performs by the *stochastic gradient descend* (SGD) algorithm, as well as increase the training speed multiple times via adapting the *batch gradient descend* (BGD) algorithm for *MapReduce* paradigm.

Acknowledgements

This work is supported by China Scholarship Council, National Program on Key Basic Research Project (973 Program) under Grant 2013CB329304, National Science Foundation of China (NSFC) under Grant No.61433018 and No.61373075, Proteus Project of NYU. The first author conducted this research while he was a joint-supervision Ph.D. student of Tsinghua University and New York University.

Author details

¹Center for Speech and Language Technology, Research Institute of Information Technology, Tsinghua University, ROOM 1-303, BLDG FIT, 100084 Beijing, China. ²Center for Speech and Language Technologies, Division of Technical Innovation and Development, Tsinghua National Laboratory for Information Science and Technology, ROOM 1-303, BLDG FIT, 100084 Beijing, China. ³Department of Computer Science and Technology, Tsinghua University, ROOM 1-303, BLDG FIT, 100084 Beijing, China. ⁴Proteus Group, 715 Broadway, Room 703, 10003 New York City, U.S.A.. ⁵Department of Computer Science, 715 Broadway, Room 703, 10003 New York City, U.S.A.. ⁶Courant Institute of Mathematical Sciences, 715 Broadway, Room 703, 10003 New York City, U.S.A..

References

1. Bordes, A., Weston, J., Collobert, R., Bengio, Y., et al.: Learning structured embeddings of knowledge bases. In: AAAI (2011). <http://www.aaai.org/ocs/index.php/AAAI/AAAI11/paper/viewPDFInterstitial/3659/3898>
2. Miller, G.A.: Wordnet: a lexical database for english. *Communications of the ACM* **38**(11), 39–41 (1995)
3. Bollacker, K., Cook, R., Tufts, P.: Freebase: A shared database of structured general human knowledge. In: AAAI, vol. 7, pp. 1962–1963 (2007). <http://www.aaai.org/Papers/AAAI/2007/AAAI07-355.pdf>
4. Bollacker, K., Evans, C., Paritosh, P., Sturge, T., Taylor, J.: Freebase: a collaboratively created graph database for structuring human knowledge. In: *Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data*, pp. 1247–1250 (2008). ACM. <http://dl.acm.org/citation.cfm?id=1376746>

5. Carlson, A., Betteridge, J., Kisiel, B., Settles, B., Jr., E.R.H., Mitchell, T.M.: Toward an architecture for never-ending language learning. In: Proceedings of the Twenty-Fourth Conference on Artificial Intelligence (AAAI 2010) (2010)
6. Bordes, A., Usunier, N., Garcia-Duran, A., Weston, J., Yakhnenko, O.: Translating embeddings for modeling multi-relational data. In: Advances in Neural Information Processing Systems, pp. 2787–2795 (2013)
7. Fan, M., Zhou, Q., Chang, E., Zheng, T.F.: Transition-based knowledge graph embedding with relational mapping properties. In: Proceedings of the 28th Pacific Asia Conference on Language, Information, and Computation, pp. 328–337. Department of Linguistics, Chulalongkorn University, Phuket, Thailand (2014). <http://www.aclweb.org/anthology/Y14-1039>
8. Fan, M., Zhou, Q., Zheng, T.F., Grishman, R.: Probabilistic belief embedding for knowledge base completion. arXiv preprint arXiv:1505.02433 (2015)
9. Fan, M., Zhou, Q., Zheng, T.F.: Learning embedding representations for knowledge inference on imperfect and incomplete repositories. arXiv preprint arXiv:1503.08155 (2015)
10. Fan, M., Cao, K., He, Y., Grishman, R.: Jointly embedding relations and mentions for knowledge population. CoRR **abs/1504.01683** (2015)
11. Fan, M., Zhou, Q., Zheng, T.F., Grishman, R.: Large margin nearest neighbor embedding for knowledge representation. CoRR **abs/1504.01684** (2015)
12. Wang, Z., Zhang, J., Feng, J., Chen, Z.: Knowledge graph and text jointly embedding. In: Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP), pp. 1591–1601. Association for Computational Linguistics, ??? (2014). <http://aclweb.org/anthology/D14-1167>
13. Wang, Z., Zhang, J., Feng, J., Chen, Z.: Knowledge graph embedding by translating on hyperplanes. In: Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence, July 27–31, 2014, Québec City, Québec, Canada., pp. 1112–1119 (2014). <http://www.aaai.org/ocs/index.php/AAAI/AAAI14/paper/view/8531>
14. Dean, J., Ghemawat, S.: Mapreduce: simplified data processing on large clusters. Communications of the ACM **51**(1), 107–113 (2008)
15. Chu, C., Kim, S.K., Lin, Y.-A., Yu, Y., Bradski, G., Ng, A.Y., Olukotun, K.: Map-reduce for machine learning on multicore. Advances in neural information processing systems **19**, 281 (2007)
16. Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space. arXiv preprint arXiv:1301.3781 (2013)